

METHOD AND APPARATUS FOR LINE CARD
REDUNDANCY IN A COMMUNICATION SWITCH

5

Technical Field of the Invention

The invention relates generally to communication systems and more particularly to a method and apparatus for line card redundancy in switches used in such communication systems.

10

Background of the Invention

Communication between a calling party and a called party may be established over a communication network. Such a communication network may include a plurality of network switches that transport data. In typical cell-based and packet-based communication networks, these switches include a switch core and a number of line cards. Each of the line cards provides a communication link to another switch in the network. The switch core allows inputs received from one line card to be provided on the output of another line card. This flexible routing of data allows for efficient communication between the various switches in the network.

15

20

In operation, line cards within a switch may experience faults or other errors that cause the line card to fail. In such instances, it is desirable to have a redundant line card that can be switched into an active mode and utilized to insure that data communication with the network is not interrupted. However, in order to provide a redundant line card for each line card in the system, prior art systems required that the effective bandwidth of the overall switch be cut in half. This is due to the fact that switch cores within the switch are only able to couple to a limited number of line cards. As such, if redundant line cards are included in the system, some of the coupling capabilities of the switch core will be consumed by these redundant line cards. As the redundant line cards are only utilized in the event of a problem with the normally active line cards, the bandwidth of the overall switch is effectively halved.

25

30

564220" 4654260

Another prior art solution includes an additional switch core in the system, where the additional switch core is coupled to a complete array of line cards. These line cards are coupled to the external data connections in the same manner as the line cards of the original switch core. When an error occurs with one of the line cards coupled to the original switch core, the entire system switches over to utilize the additional switch core provided in the system. However, this transition effectively eliminates any future redundancy for all of the line cards in the system. In addition to this, including an additional switch core within the switch adds to the cost and complexity of the switch.

Therefore, a need exists for a method and apparatus for providing line card redundancy in a communication switch without reducing the effective bandwidth of the switch.

Brief Description of the Drawings

Figure 1 illustrates a prior art communication switch;

Figure 2 illustrates a block diagram of a communication switch in accordance with the present invention;

Figure 3 illustrates a block diagram of a line card manager in accordance with the present invention;

Figure 4 illustrates a block diagram of a cell that may be used in conjunction with the switch of Figure 2;

Figure 5 illustrates a block diagram of another line card manager in accordance with the present invention; and

Figure 6 illustrates a flow diagram of a method for managing line cards in accordance with the present invention.

Detailed Description of a Preferred Embodiment

Generally, the present invention provides a method and apparatus for line card redundancy in a communication switch. An arbiter is coupled to the input of the switch core where a single line card would be coupled in a prior art solution. The arbiter in turn is coupled to a line card pair where a first line card of the line card pair is considered an active line card, whereas the second line card of the line card pair is considered the

redundant, or inactive, line card. The active and inactive line cards are determined based on selection information received by the arbiter. Based on the selection information, the arbiter preferentially provides ingress data from the active line card to the corresponding input of the switch core. A router is coupled to the line card pair and to an output of the switch core. The router passes egress data from the switch core to at least one of the first and second line cards based on routing information included in the egress data. By enabling an input and output of the switch core to couple to more than one line card, redundancy can be established without removing effective bandwidth from the communication switch. If a failure on an active line card or other error occurs, selection information can be provided to the arbiter to indicate that it should pass ingress data from the newly activated, formerly redundant line card.

The invention can be better understood with reference to Figures 1-6. Figure 1 illustrates a switch core 10 operably coupled to a plurality of line cards 20-35. Each of the line cards provides ingress data to the switch core 10 and receives egress data from the switch core 10. In this prior art solution, the redundancy only exists if only a portion of the line cards is used as active line cards at any one time. Thus, if line card 20 is designated an active line card, and fails for some reason, line card 21 may be utilized as a replacement for the defective line card. Although this may allow the switch to continue operating, it effectively halves the bandwidth of the switch in operation.

Figure 2 illustrates a switch that provides redundancy yet does not compromise any bandwidth. The system includes the switch core 10, a plurality of line card managers 50-65, a plurality of active line cards 70-85, and a plurality of initially redundant line cards 90-105. Each of the line card managers 50-65 allows the switch core 10 to couple to a pair of line cards. One line card of the line card pair is designated as an active line card, which may be the line cards 70-85. The active line card will be utilized until some fault or error causes the line card manager to switch from the active line card to the formerly inactive line card. Thus, the line cards 90-105 are initially set up as inactive line cards and are provided for redundancy.

The switch core 10 of Figure 2 is preferably a NxN switch core that has N inputs and N outputs. The switch core 10 passes data received from the inputs to the outputs based on routing tags. More preferably, the switch core 10 is a 16x16 switch core that,

through the use of the line card managers 50-65, is able to couple to 32 line cards. For other values of N, the number of line cards to which the switch core is able to couple is preferably equal to $(2 * N)$. The redundant line card switch illustrated in Figure 2 may be utilized in either a cell-based or a packet-based network, and most preferably is utilized in an asynchronous transfer mode (ATM) network.

Figure 3 illustrates a more detailed view of the line card manager 50 coupled to line cards 70 and 90. The line card manager 50 includes an arbiter 110, a router 120, and preferably includes a buffer 130 and filters 140 and 142. The arbiter 110 is operably coupled to the first line card 70 and the second line card 90. The arbiter 110 is also coupled to an input 170 of the switch core. The arbiter provides ingress data received from one of the first and second line cards to the corresponding input of the switch core based on selection information 160. Preferably, the selection information 160 determines an active line card from the first and second line cards 70 and 90. The arbiter will preferentially pass the data from the active line card before it passes inactive line card data. More preferably, the arbiter 110 will only pass inactive line card data when idle cells or packets are included in the active line card data.

Although the amount of data passed by the inactive line card is preferably minimal, enough data may be permitted to pass from the inactive line card to perform diagnostic and other low bandwidth applications. In order to insure that some inactive line card data is passed through to the switch core input 170, the active line card may intentionally insert idle cells into its data stream.

A buffer 130 may be included in the system upstream from the arbiter 110 such that data cells from the line cards 70 and 90 can be held until it is possible for the arbiter 110 to pass them to the switch core input 170. Preferably, the buffer 130 includes a first-in/first-out (FIFO) buffer corresponding to each of the line cards. The FIFO 132 will buffer ingress data received from the line card 70, whereas the FIFO 134 will buffer ingress data received from the line card 90. The size or capacity of the FIFOs may be determined based on whether the line card to which the FIFO corresponds is normally utilized as the active or inactive line card. If the line card is typically inactive, the FIFO may be smaller than that of the active line card. More preferably, the FIFO corresponding to the inactive line card may be sized based on the number of idle cells

typically passed by the active line card, as this determines how often data from the inactive FIFO will be passed to the switch core.

Filters 140 and 142 may also be included in the ingress data path of the line cards 70 and 90. The filters 140 and 142 are preferably configured by a filter configuration block 150. The filters 140 and 142 may be used to insure that only relevant data from the line cards 70 and 90 is presented to the buffer 130. More preferably, the filters 140 and 142 are configured based on their coupling to an active or inactive line card such that a filter corresponding to an active line card passes data that is expected to be passed through to the switch core input 170. A filter corresponding to an inactive line card will most likely filter out the majority of data, only passing data relevant to diagnostic or other low bandwidth applications. In order to configure the filters in such a manner, registers may be included in the filter configuration block 150 that are selected and used to configure the filters 140 and 142 based on whether they couple to an active or inactive line card. Thus, the selection information 160 may be routed to the filter configuration block 150 such that it selects either an active setting register 150 or an inactive setting register 154 to configure a particular filter.

The determination of whether or not a filter will pass a certain data cell or packet may be based on the data type of the received data cell. The data type may indicate whether the data would be relevant to an active line card or an inactive line card.

Referring to Figure 4, the data type is preferably determined based on a type field included in the header of a data cell or packet. Figure 4 shows a data cell 200 that includes a cell header 210 and a payload 202. Included in the cell header 210 is a type field 212 that may be used by the filters 140 and 142 to determine whether or not the data cell should be passed through to the buffer 130. Thus, the active and inactive settings 152 and 154 configure the filters 140 and 142 such that they pass selected data types and reject other data types.

Returning to Figure 3, egress data from the switch core output 180 is directed by the router 120 to one or more of the line cards 70 and 90. Preferably, the router 120 will redirect the egress data to both of the line cards 70 and 90 such that if a problem does occur, the transition from one line card to the other will be smooth. In order to determine to which of the line cards 70 and 90 to provide the egress data, the router 120 examines

routing information included in the egress data. Preferably, a bit is included in the header of the egress data cells that corresponds to each of the line cards, and when a bit is set to a valid state, the router 120 will provide the egress data to the particular line card corresponding to that bit.

5 Returning to Figure 4, it can be seen that the cell header 210 includes a first bit 214 and a second bit 216. In one embodiment, the first bit 214 corresponds to the line card 70, and the second bit 216 corresponds to the line card 90. When the first bit is set to a valid state, egress data will be passed to the line card 70, and when the second bit 216 is set to a valid state, egress data will be passed to the line card 90. This enables data
10 to be selectively passed to one or both of the line cards 70 and 90.

 In another embodiment, the first bit 214 is known to correspond to the active line card, whereas the second bit 216 is known to correspond to the inactive line card. Thus, when the first bit 214 is set, the egress data will be passed to the active line card, which may be line card 70 or 90 based on the selection information 160. Similarly, the second
15 bit 216 will correspond to the inactive line card as determined by the selection information 160. Thus, if the first bit in this embodiment were set to a valid state, the data would be passed to the active line card and if the second bit 216 is set, the data is passed to the inactive line card.

 It should be noted that in both embodiments that use the two-bit routing
20 information to control distribution of the egress data, data can be provided to either one of the line cards 70 or 90 or to both of the line cards 70 and 90 based on the settings of the two bits. This is preferable to utilizing a single bit to determine which of the line cards is provided with the egress data, where one bit state corresponds to one line card and the other bit state corresponds to the other line card. Such an embodiment would
25 only allow the egress data to be passed to one or the other of the line cards, which may not be desirable.

 Although the preferred embodiment of Figures 2 and 3 illustrates one redundant line card for each active line card in the system, in order to conserve resources, a single line card may be used as the redundant line card for more than one active line card.

30 Figure 5 illustrates an embodiment that includes three line cards 570, 580, and 590 that are coupled to two switch core inputs 670 and two switch core outputs 680. The line card

manager 550 is used control the flow of ingress and egress data between the line cards 570, 580, and 590 and the switch core inputs 670 and outputs 680. In the example illustrated, the line card 570 is preferably the active line card coupled to the input [0] of the switch core. Similarly, the line card 590 is preferably the active line card coupled to the switch core input [1]. These line cards are also coupled to receive the corresponding egress data from the respective switch core outputs 680.

The line card 580 is included in the system as a redundant line card which is shared by the line cards 570 and 590. In order to utilize the line card 580 as a shared redundant resource, an arbiter 610 is included that is operably coupled to all three of the line cards 570, 580, and 590. The arbiter 610 determines which of the line cards are active based on selection information 660 and routes ingress data from these active line cards to the corresponding switch core inputs 670. Similarly, the router 620 determines which of the line cards 570, 580, and 590 are to receive egress data corresponding to two of the outputs from the switch core. In the embodiment illustrated, a two-bit quantity may still be included as routing information in the header, wherein the two-bit quantity selects between the active line card for an output and the shared redundant line card. Data cells corresponding to one of the active line cards may be provided to the shared line card, such that if that particular active line card fails, the transition will be smooth. In the case where the other active line card fails, the changeover to the redundant line card may require some type of initialization or start-up operations. Alternatively, a three-bit quantity may be included as routing information in the header of a cell such that the router 620 may distribute a particular data cell to any one or more of the three line cards based on this routing information.

As before, the ingress data path preferably includes a buffer 130 that may include individual FIFOs 632, 634, and 636 corresponding to each of the ingress data paths. Similarly, filters 640, 642, and 644 may be included in the ingress data paths such that ingress data is passed based on its data type as determined by the type field in the cell header. As before, these filters may be configured using a filter configuration block 650 that includes active and inactive settings that are preferably stored in registers 652 and 654.

It should be noted that the ratio of active line cards to redundant line cards should be based on the particular network in which the line cards are utilized, and the example illustrated in Figure 5 is merely for illustration purposes. Thus, three active line cards may share a single redundant line card, or five line cards may share three redundant line cards, etc. Similarly, in the embodiment illustrated in Figure 2, all of the line cards in the system may not include a redundant line card, such that only a portion of the line cards in a particular switch have a potential for a redundant back-up.

Figure 6 illustrates a flow diagram of a method for managing line cards in a system that includes redundant line cards. The method begins at step 300 where data is received from a first line card and a second line card. As was described with respect to Figures 3 and 5, the ingress data may be filtered at step 302 based on its data type. At step 304, the ingress data is buffered. Preferably, this is accomplished using FIFOs as was described earlier. The filtering performed at step 302 can insure that only relevant inactive line card data is passed through to the inactive line card buffer, which may be of limited size. This enables diagnostics information of a certain data type to be passed through and buffered without being overwhelmed and pushed out of the limited sized buffer by other data types.

At step 306, ingress data is selected from the received data based on an active select signal. The active select signal determines an active line card and an inactive line card from the first and second line cards. Preferably, the ingress data is selected based on step 308, such that active line card data is preferentially selected over inactive line card data. Thus, as was described earlier, the line card data corresponding to the active line card will be passed unless idle data cells are included in the data stream, in which case inactive line card data may be allowed to trickle through. This small amount of bandwidth provided to the inactive line card may allow for diagnostics or other low bandwidth applications to be maintained.

At step 310, the ingress data is provided to an input of the switch core. As before, the switch core preferably includes a plurality of inputs and a plurality of outputs where data received on the inputs is passed to any one of the outputs based on routing tags included in the data cells.

03275934.03249
65420.465260

At step 314, egress data is received from an output of the switch core. At step 314, this egress data is selectively provided to at least one of the first and second line cards based on routing information included in the egress data. Preferably, the routing information included in the egress data is a two-bit quantity, where one bit corresponds to one of the line cards and the bit corresponds to the other line card. In one embodiment, the first bit will correspond to a particular line card, whereas in another embodiment, the first bit will correspond to the active line card. Similarly, the second bit will either correspond to another particular line card or to the inactive line card. Note that the ordering of these two bits is arbitrary. When a bit is set to the active state, the egress data will be passed to the line card to which it corresponds.

The present invention overcomes the limitations of the prior art, which forced a reduction in effective bandwidth of a switch when redundant line cards were included. In a prior art implementation that included 16 line cards coupled to a 16x16 switch core, only eight of the line cards could be active, and the other eight were regulated to redundant status. This is inefficient use of bandwidth in the system, and increases costs in packet- and cell-based networks.

The present invention allows for a means to couple 32 full-capacity line cards to a single 16x16 switch core. The ingress data to a particular input of the 16x16 switch core is selected from two line card inputs, each of which includes a FIFO for buffering purposes. An active line card is chosen from the two line cards, and an arbiter is included which preferentially passes data from the active line card. Preferably, the arbiter only passes inactive line card data when idle data cells are included in the active line card data stream.

Thus, the present invention allows for line card redundancy without compromising overall bandwidth of a switch. It should be understood that the implementation of other variations and modifications of the invention in its various aspects should be apparent to those of ordinary skill, and that the invention is not limited to the specific embodiments described. It is therefore contemplated to cover by the present invention, any and all modifications, variations, or equivalents that fall within the spirit and scope of the basic underlying principles disclosed and claimed herein.